

NETWORK ROUTE CONTROL

Field of the Invention

The present invention generally relates to the computer networks and computer systems and, more specifically, to implementation of route control for networks that are

5 multi-homed to multiple network providers.

Background of the Invention

Large enterprises and content providers, who depend on the Internet to operate their businesses, require a high level of reliability from their network connections.

Increasingly, these large consumers and producers of network data are turning to multi-

10 homing as a technique to achieve resilience to service interruptions. Multi-homing is defined simply as a customer (or Internet Service Provider ["ISP"]) network having more than one external link, either to a single ISP, or to different providers. The customer typically has its own public Autonomous System ("AS") number, and advertises its address prefixes via all of its upstream providers using Border Gateway Protocol

15 ("BGP"). See BGP4: Inter-Domain Routing in the Internet by John W. Stewart, Addison-Wesley Pub. Co. (1st ed. 1998) or Internet Engineering Task Force (IETF) Request for Comments (RFC) 1771, for a general discussion of the Border Gateway Protocol.

While multi-homing to multiple providers is motivated primarily by a need for link-level and provider-level fault tolerance, recently developed "intelligent route control" devices and services allow subscribers to leverage multi-homing for more than just increased resilience. For example, performance to different parts of the network may vary depending on which upstream provider is used. In such situations, careful route selection can significantly improve performance. Even availability can be managed to some extent by choosing ISPs that have sufficiently diverse connectivity to destinations of interest.

Route control solutions require two main components: measurement of the candidate links, and control to steer outgoing traffic over the best performing link. The measurement component typically measures the delay and loss over upstream provider links to various destinations in the network. Using these measurements, the best provider to reach a particular destination can be identified. Once the best provider is identified, however, traffic to the destination must be directed over the appropriate link. Current solutions achieve this link/ISP selection by interacting with the BGP router(s) connecting the data center to the ISP, usually with an external device adjunct to the router making this selection.

Current solutions are implemented as network appliances deployed in data center or enterprise networks. These appliances typically measure delays and availability over each ISP link using passive and active probing techniques such as observing Transmission Control Protocol ("TCP") connection establishment delay, or using network pings. These low-level measurements are not application-specific, which means that the delay measurements must be translated into a metric that is meaningful for the application communicating over the network. In the currently available solutions, route control is done by modifying BGP forwarding tables so that traffic destined for a particular network uses the best-performing ISP based on observed performance. This requires installation of the appliance near the edge routers in the network, and complex router configuration to allow the appliance to direct traffic over links that may be contrary to the choice made by the standard BGP protocol.

Summary of the Invention

The invention broadly contemplates any network-connected general purpose computer performing route control functions, without the need for external network appliances or direct interaction with BGP routers, as done in traditional approaches.

In summary, one aspect of the invention provides a method for network route control, the method comprising the steps of: establishing a connection between a general

purpose computer and arrangements for linking to multiple internet service providers (ISPs); measuring relevant performance and availability metrics of said links; and said computer directs network traffic to the best link based upon said relevant performance and availability metrics.

5 Another aspect of the present invention provides an apparatus permitting a general purpose computer to perform route control, the apparatus comprising: an arrangement for establishing a link between a general purpose computer and multiple internet service providers (ISPs); an arrangement allowing said computer to measure at least one relevant performance metric of said links; and an arrangement allowing said computer to select the
10 best route based upon said link performance.

 Furthermore, an additional aspect of the invention provides a program storage device readable by machine, tangibly embodying a program of instructions executable by the machine to perform method steps for general purpose computer route control, said method comprising the steps of: establishing a connection between a general purpose
15 computer and arrangements for linking to multiple internet service providers (ISPs); measuring relevant performance metrics of said links; and said computer directs network traffic to the best link based upon said relevant performance and availability metrics.

For a better understanding of the present invention, together with other and further features and advantages thereof, reference is made to the following description, taken in conjunction with the accompanying drawings, and the scope of the invention that will be pointed out in the appended claims.

5 **Brief Description of the Drawings**

Fig. 1 is a schematic diagram showing the configuration of the computer network system in which network subscribers are multi-homed to multiple network providers or Internet service providers using a route control device in accordance with the prior art.

Fig. 2. is a schematic diagram showing the configuration of the computer network
10 system using a general purpose computer to select Internet routes according to one embodiment of the present invention.

Fig. 3. is a schematic diagram showing the configuration of the computer network system using a general purpose computer to select Internet routes according to another embodiment of the present invention.

15 **Description of the Preferred Embodiments**

The present invention applies to environments in which network subscribers are multi-homed to multiple network providers, or ISPs. An example of this scenario is

shown in Fig.1, a diagram showing the configuration of the computer network system in which network subscribers are multi-homed to multiple network providers or Internet service providers using a route control device in accordance with the prior art.

In Fig. 1, the enterprise 100 is connected to two ISPs, shown as "ISP 1" 160 and "ISP 2" 170, with edge router 130 on the customer premises connected to border routers 140 and 150 in each of the ISPs. In accordance with present practice, route control is achieved by having a route control device 120 monitor the delay or loss on each of the ISP links using active or passive measurements, and then updating the BGP forwarding tables on the edge router to reflect the route control decision for destinations of interest.

The present invention removes the need for a dedicated route control device and instead permits general purpose computers to perform route control functions themselves. In the preferred embodiment, a general purpose computer hosting applications which can benefit from route control supports the Multi-protocol Label Switching (MPLS) protocol in its native operating system networking stack. For a general description of MPLS, see Internet Engineering Task Force (IETF) Request for Comments (RFC) 3031 . Moreover, the devices in the customer premises network (e.g., switches, routers) between the general purpose computer and the edge routers are MPLS-enabled. This implies that they can perform the label switching function of Label Switched Routers (LSRs) (this function is

also described in IETF RFC 3031). The network administrator, through some manual or automated means, establishes multiple MPLS label-switched paths (LSPs) between the general purpose computer and the edge routers, with one LSP for each service provider link. At the edge router, the label corresponding to each LSP is used to determine on which link the packet is forwarded. This scenario is shown in Fig. 2, a diagram showing the configuration of the computer network system using a general purpose computer to select Internet routes according to one embodiment of the present invention.

In Fig. 2, utilizing router 230, LSP 1 at 210 and LSP2 at 220 are configured between the general purpose computer 200 and the edge router 130. Packets sent by the general purpose computer with a label corresponding to LSP 1 is forwarded by the edge router on the link to ISP 1. Packets sent on LSP 2 are similarly forwarded toward ISP 2.

For the general purpose computer to make route control decisions, it must be able to measure the relative performance or availability on each of the provider links. This can be done with either active or passive measurement of metrics of interest. For active probing, the general purpose computer can send multiple probes toward a destination, one on each LSP to measure the performance on each link. For example Internet Control Message Protocol ("ICMP") echo (i.e., ping) packets can be sent over each link to compare the round-trip delay or packet loss on each link. For a general description of this

process, see chapters 6 and 7 of TCP/IP Illustrated Volume 1: The Protocols by W. Richard Stevens, Addison-Wesley Pub. Co. (1st ed. 1994).

For passive performance measurement of each ISP link, the applications running on the general purpose computer can respond to incoming requests using a specified link by sending the reply on different LSPs, and thus different ISP links. By using normal application interactions with actual clients, the general purpose computer has the flexibility to measure performance metrics in an application-specific manner.

Once sufficient performance or availability measurement samples are obtained for each ISP link, the general purpose computer can direct outgoing application traffic over the most appropriate (e.g., best performing) link. The route control decision is implemented in a manner similar to that used for collecting measurements, namely by sending packets on the appropriate LSP. The route control decision may be based on a number of metrics, including, but not limited to, network delay, loss, and throughput or application-layer response time. The route control decision can be made on a per-destination basis, or on using some level of aggregation to group clients.

In another embodiment, the general purpose computer can use different Virtual Local Area Network ("VLAN") identifiers within the subscriber network to correspond to different outgoing links. VLANs are generally described in The Switch Book by Rich

Seifert, John Wiley and Sons (1st. ed. 2000). Referring now to Fig. 3, the edge 130 uses
VLAN identifiers 310 and 320 to forward packets on a specific IP link. By transmitting
packets on a specified VLAN, the general purpose computer 200 can ensure that a
particular ISP link is used. These VLANs are established within the subscriber network
5 using some manual or automated means. The measurements and route control
mechanism operate similarly to the MPLS embodiment described above. In this method,
instead of using MPLS labels to designate the outgoing link, VLAN identifiers are used
for this purpose.

Another embodiment uses Internet Protocol ("IP") level tunnels (such as IP-IP, or
10 Generic Routing Encapsulation (GRE) based) to direct packets to a specified outgoing
ISP link (see IETF RFC 2003 and RFC 2784). These tunnels, once established, can be
assigned to different virtual interfaces on the general purpose computer, with each
interface corresponding to a different ISP link. Packets can then be directed by the
general purpose computer by transmitting them on a specified virtual interface. At the
15 edge router, packets that arrive on a specified tunnel are forwarded on the corresponding
ISP link.

It is to be understood that the present invention, in accordance with at least one
presently preferred embodiment, includes an arrangement for establishing a link between

a general purpose computer and multiple internet service providers (ISPs), an arrangement allowing said computer to measure relevant performance metric of said links, and an arrangement allowing said computer to select the best route based upon said link performance, all of which may be implemented on at least one general-purpose computer running suitable software programs. These may also be implemented on at least one Integrated Circuit or part of at least one Integrated Circuit. Thus, it is to be understood that the invention may be implemented in hardware, software, or a combination of both.

If not otherwise stated herein, it is to be assumed that all patents, patent applications, patent publications and other publications (including web-based publications) mentioned and cited herein are hereby fully incorporated by reference herein as if set forth in their entirety.

Although illustrative embodiments of the present invention have been described herein with reference to the accompanying drawings, it is to be understood that the invention is not limited to those precise embodiments, and that various other changes and modifications may be affected therein by one skilled in the art without departing from the scope or spirit of the invention.